

Workgroup: BFD Working Group

Internet-Draft:

draft-ietf-bfd-unaffiliated-echo-10

Updates: [5880](#) (if approved)

Published: 28 September 2023

Intended Status: Standards Track

Expires: 31 March 2024

Authors: W. Cheng	R. Wang	X. Min, Ed.
China Mobile	China Mobile	ZTE Corp.
R. Rahman	R. Boddireddy	
Equinix	Juniper Networks	

Unaffiliated BFD Echo

Abstract

Bidirectional Forwarding Detection (BFD) is a fault detection protocol that can quickly determine a communication failure between two forwarding engines. This document proposes a use of the BFD Echo where the local system supports BFD but the neighboring system does not support BFD. BFD Control packet and its processing procedures can be executed over the BFD Echo port where the neighboring system only loops packets back to the local system.

This document updates RFC 5880.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 31 March 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents

(<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
 - [1.1. Conventions Used in This Document](#)
- [2. Unaffiliated BFD Echo Procedures](#)
- [3. Updates to RFC 5880](#)
- [4. Unaffiliated BFD Echo Applicability](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
- [7. Acknowledgements](#)
- [8. Contributors](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

To minimize the impact of device/link faults on services and improve network availability, in the single-hop cases a network device needs to be able to quickly detect faults in communication with adjacent devices. Measures can then be taken to promptly rectify the faults to ensure service continuity.

BFD [[RFC5880](#)] is a low-overhead, short-duration method to detect faults on the communication path between adjacent forwarding engines. The faults can be on interfaces, data link(s), and even the forwarding engines. It is a single, unified mechanism to monitor any media and protocol layers in real time.

BFD defines Asynchronous and Demand modes to satisfy various deployment scenarios. It also supports an Echo function to reduce the device requirement for BFD. When the Echo function is activated, the local system sends BFD Echo packets and the remote system loops back the received Echo packets through the forwarding path. If several consecutive BFD Echo packets are not received by the local system, then the BFD session is declared to be Down.

When using BFD Echo function, there are two typical scenarios as below:

- *Full BFD protocol capability with affiliated Echo function. This scenario requires both the local device and the neighboring device to support the full BFD protocol.

- *BFD Echo-Only method without full BFD protocol capability. This scenario requires only the local device to support sending and demultiplexing BFD Control packets. In this scenario, the BFD Control packets are sent over the BFD Echo port, but that the processing procedures for Asynchronous mode are used with the modifications described in this document. Note that this method monitors the connectivity to a system over a specific interface and does not verify the availability of a specific IP address at that system.

The former scenario is referred to as affiliated BFD Echo, which is not changed by this document in any way. The latter scenario is referred to as Unaffiliated BFD Echo, which is specified in this document.

Section 5 of [[RFC5880](#)] indicates that the payload of an affiliated BFD Echo packet is a local matter and hence its contents are outside the scope of that specification. This document, on the other hand, specifies the contents of the Unaffiliated BFD Echo packet and what to do with them.

Section 6.2.2 of [[BBF-TR-146](#)] describes one use case of the Unaffiliated BFD Echo. Section 2 of [[I-D.wang-bfd-one-arm-use-case](#)] describes another use case of the Unaffiliated BFD Echo.

This document describes the use of the Unaffiliated BFD Echo over IPv4 and IPv6 for single IP hop.

1.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2. Unaffiliated BFD Echo Procedures



Figure 1: Unaffiliated BFD Echo diagram

As shown in Figure 1, device A supports BFD, whereas device B does not support BFD. Device A would send Unaffiliated BFD Echo packets, and after receiving the Unaffiliated BFD Echo packets sent from device A, the one-hop-away BFD peer device B immediately loops them back by normal IP forwarding, this allows device A to rapidly detect a connectivity loss to device B. Note that device B would not intercept any received Unaffiliated BFD Echo packet or parse any BFD protocol field within the Unaffiliated BFD Echo packet.

For unaffiliated echo, a Unaffiliated BFD Echo session is created on device A, and the Unaffiliated BFD Echo session MUST follow the BFD state machine defined in Section 6.2 of [[RFC5880](#)], except that the received state is not sent but looped back from the remote system. Unaffiliated BFD Echo does not use the AdminDown state. BFD Control packets are transmitted and received as Unaffiliated BFD Echo packets using destination UDP port 3785, as defined in [[RFC5881](#)]. The procedures for BFD Async sessions are executed for the looped BFD Control packets as per [[RFC5880](#)], including validation and authentication.

Once a Unaffiliated BFD Echo session is created on device A, it starts sending Unaffiliated BFD Echo packets. Unaffiliated BFD Echo packets with zeroed "Your Discriminator" are demultiplexed to the proper session based on the source IP address or UDP source port, once the remote system loops back the local discriminator, all further received packets are demultiplexed based on the "Your Discriminator" field only, which is conformed to the procedure specified in Section 6.3 of [[RFC5880](#)]. An Unaffiliated BFD Echo packet follows the same encapsulation rules as for a BFD Echo packet as specified in Section 4 of [[RFC5881](#)]. All Unaffiliated BFD Echo packets for the session MUST be sent with a Time to Live (TTL) or Hop Limit value of 255, and received with a TTL or Hop Limit value of 254, otherwise the received packets MUST be dropped [[RFC5082](#)].

Within the Unaffiliated BFD Echo packet, the "Desired Min TX Interval" and "Required Min RX Interval" defined in [[RFC5880](#)] SHOULD be populated with a certain value, which can avoid unset value being a potential vector for disclosure of uninitialized memory. A

suggested value is 1 second (1,000,000 microseconds). These values, however, MUST be ignored on receipt. Furthermore, these values MUST NOT be used to calculate the Detection Time.

The "Required Min Echo RX Interval" defined in [\[RFC5880\]](#) SHOULD be populated with a certain value. A suggested value is 0. This value MUST be ignored on receipt. The transmission interval for Unaffiliated BFD Echo packets in the Up state MUST be provisioned on device A. The Unaffiliated BFD Echo feature depends on device B performing IP forwarding (actually IP redirect) functionality. While such functionality may normally be expected to be supported on a router, it may not be enabled on a host by default. The method for provisioning device B to loop back Unaffiliated BFD Echo packets is outside the scope of this document.

Similar to what's specified in [\[RFC5880\]](#), the Unaffiliated BFD Echo session begins with the periodic, slow transmission of Unaffiliated BFD Echo packets. The slow transmission rate SHOULD be no less than one second per packet, until the session is Up. After the session is Up, the provisioned transmission interval is used. When the Unaffiliated BFD Echo session goes Down, the slow transmission rate is resumed. The "Detect Mult" defined in [\[RFC5880\]](#) MUST be set to a value provisioned on device A. When the bfd.SessionState is Up and a Detect Mult number of Unaffiliated BFD Echo packets have not arrived at device A as they should, the device A MUST set bfd.SessionState to Down and bfd.LocalDiag to 2 (Echo Function Failed), as specified in Section 6.8.5 of [\[RFC5880\]](#).

In summary, the Unaffiliated BFD Echo packet reuses the format of the BFD Control packet defined in [\[RFC5880\]](#), and the fields within the Unaffiliated BFD Echo packet are populated as follows:

- *My Discriminator MUST be set to the provisioned local discriminator.

- *Your Discriminator MUST be set to 0 initially, and then MUST be set to the same as My Discriminator looped back.

- *Desired Min TX Interval SHOULD be set to a certain value. A suggested value is 1 second (1,000,000 microseconds).

- *Required Min RX Interval SHOULD be set to a certain value. A suggested value is 1 second (1,000,000 microseconds).

- *Required Min Echo RX Interval SHOULD be set to a certain value. A suggested value is 0.

- *Detect Mult MUST be set to the provisioned maximum allowable number of consecutively lost Unaffiliated BFD Echo packets.

3. Updates to RFC 5880

The Unaffiliated BFD Echo described in this document reuses the BFD Echo function as described in [[RFC5880](#)] and [[RFC5881](#)], but does not require BFD Asynchronous or Demand mode. When using the Unaffiliated BFD Echo, only the local system has the BFD protocol enabled; the remote system just loops back the received BFD Echo packets as regular data packets.

This document updates [[RFC5880](#)] with respect to its descriptions on the BFD Echo function as follows.

The 4th paragraph of Section 3.2 of [[RFC5880](#)] is updated as below:

OLD TEXT

An adjunct to both modes is the Echo function.

NEW TEXT

An adjunct to both modes is the Echo function, which can also be running independently.

OLD TEXT

Since the Echo function is handling the task of detection, the rate of periodic transmission of Control packets may be reduced (in the case of Asynchronous mode) or eliminated completely (in the case of Demand mode).

NEW TEXT

Since the Echo function is handling the task of detection, the rate of periodic transmission of Control packets may be reduced (in the case of Asynchronous mode) or eliminated completely (in the case of Demand mode). The Echo function may also be used independently, with neither Asynchronous nor Demand mode.

The 3rd and 9th paragraphs of Section 6.1 of [[RFC5880](#)] are updated as below:

OLD TEXT

Once the BFD session is Up, a system can choose to start the Echo function if it desires and the other system signals that it will allow it. The rate of transmission of Control packets is typically kept low when the Echo function is active.

NEW TEXT

When a system is running with Asynchronous or Demand mode, once the BFD session is Up, it can choose to start the Echo function if it desires and the other system signals that it will allow it. The rate of transmission of Control packets is typically kept low for Asynchronous mode or eliminated completely for Demand mode when the Echo function is active.

OLD TEXT

If the session goes Down, the transmission of Echo packets (if any) ceases, and the transmission of Control packets goes back to the slow rate.

NEW TEXT

In Asynchronous mode, if the session goes Down, the transmission of Echo packets (if any) ceases, and the transmission of Control packets goes back to the slow rate. Demand mode MUST NOT be active if the session goes Down.

The 2nd paragraph of Section 6.4 of [[RFC5880](#)] is updated as below:

OLD TEXT

When a system is using the Echo function, it is advantageous to choose a sedate reception rate for Control packets, since liveness detection is being handled by the Echo packets. This can be controlled by manipulating the Required Min RX Interval field (see section 6.8.3).

NEW TEXT

When a system is using the Echo function with Asynchronous mode, it is advantageous to choose a sedate reception rate for Control packets, since liveness detection is being handled by the Echo packets. This can be controlled by manipulating the Required Min RX Interval field (see section 6.8.3). Note that a system operating in Demand mode would direct the remote system to cease the periodic transmission of BFD Control packets, by setting the Demand (D) bit in its BFD Control packets.

The 2nd paragraph of Section 6.8 of [[RFC5880](#)] is updated as below:

OLD TEXT

When a system is said to have "the Echo function active" it means that the system is sending BFD Echo packets, implying that the session is Up and the other system has signaled its willingness to loop back Echo packets.

NEW TEXT

When a system in Asynchronous or Demand mode is said to have "the Echo function active" it means that the system is sending BFD Echo packets, implying that the session is Up and the other system has signaled its willingness to loop back Echo packets.

The 7th paragraph of Section 6.8.3 of [[RFC5880](#)] is updated as below:

OLD TEXT

When the Echo function is active, a system SHOULD set `bfd.RequiredMinRxInterval` to a value of not less than one second (1,000,000 microseconds). This is intended to keep received BFD

Control traffic at a negligible level, since the actual detection function is being performed using BFD Echo packets.

NEW TEXT

When the Echo function is active with Asynchronous mode, a system SHOULD set `bfd.RequiredMinRxInterval` to a value of not less than one second (1,000,000 microseconds). This is intended to keep received BFD Control traffic at a negligible level, since the actual detection function is being performed using BFD Echo packets. While a system operating in Demand mode would not receive BFD Control traffic.

The 1st and 2nd paragraphs of Section 6.8.9 of [[RFC5880](#)] are updated as below:

OLD TEXT

BFD Echo packets MUST NOT be transmitted when `bfd.SessionState` is not Up. BFD Echo packets MUST NOT be transmitted unless the last BFD Control packet received from the remote system contains a nonzero value in Required Min Echo RX Interval.

NEW TEXT

When a system is using the Echo function with either Asynchronous or Demand mode, BFD Echo packets MUST NOT be transmitted when `bfd.SessionState` is not Up, and BFD Echo packets MUST NOT be transmitted unless the last BFD Control packet received from the remote system contains a nonzero value in Required Min Echo RX Interval.

OLD TEXT

BFD Echo packets MAY be transmitted when `bfd.SessionState` is Up. The interval between transmitted BFD Echo packets MUST NOT be less than the value advertised by the remote system in Required Min Echo RX Interval...

NEW TEXT

When a system is using the Echo function with either Asynchronous or Demand mode, BFD Echo packets MAY be transmitted when `bfd.SessionState` is Up, and the interval between transmitted BFD Echo packets MUST NOT be less than the value advertised by the remote system in Required Min Echo RX Interval...

4. Unaffiliated BFD Echo Applicability

Some devices that would benefit from the use of BFD may be unable to support the full BFD protocol. Examples of such devices include servers running virtual machines, or Internet of Things (IoT) devices.

Unaffiliated BFD Echo can be used when two devices are connected and only one of them supports the BFD protocol, and the other is capable of looping Unaffiliated BFD Echo packets.

5. Security Considerations

All Security Considerations from [[RFC5880](#)] and [[RFC5881](#)] apply.

Unaffiliated BFD Echo requires the remote device to loop Unaffiliated BFD Echo packets. In order to provide this service, the remote device cannot make use of Unicast Reverse Path Forwarding (URPF) [[RFC3704](#)] [[RFC8704](#)] in strict mode.

As specified in Section 5 of [[RFC5880](#)], since BFD Echo packets may be spoofed, some form of authentication SHOULD be included. Considering the Unaffiliated BFD Echo packets in this document are also BFD Control packets, the "Authentication Section" as defined in [[RFC5880](#)] for BFD Control packet is RECOMMENDED to be included within the Unaffiliated BFD Echo packet.

In order to mitigate the potential reflector attack by the remote attackers, or infinite loop of the Unaffiliated BFD Echo packets, it's RECOMMENDED to put two requirements, also known as Generalized TTL Security Mechanism (GTSM) [[RFC5082](#)], on the device looping Unaffiliated BFD Echo packets, the first one is that a packet SHOULD NOT be looped unless it has a TTL or Hop Limit value of 255, and the second one is that a packet being looped MUST NOT reset the TTL or Hop Limit value to 255, and MUST use a TTL or Hop Limit value of 254.

As stated in Section 2, in order to avoid unset values being a potential vector for disclosure of uninitialized memory, all fields of the Unaffiliated BFD Echo packet SHOULD be populated with a certain value, even if some of the fields are ignored on receipt.

6. IANA Considerations

This document has no IANA action requested.

7. Acknowledgements

The authors would like to acknowledge Ketan Talaulikar, Greg Mirsky, Santosh Pallagatti, and Aijun Wang for their careful review and very helpful comments.

The authors would like to acknowledge Jeff Haas for his guidance, insightful review and very helpful comments.

The authors would like to acknowledge Detao Zhao for the very helpful discussion.

8. Contributors

Liu Aihua
ZTE
Email: liu.aihua@zte.com.cn

Qian Xin
ZTE
Email: qian.xin2@zte.com.cn

Zhao Yanhua
ZTE
Email: zhao.yanhua3@zte.com.cn

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.
- [RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, DOI 10.17487/RFC5881, June 2010, <<https://www.rfc-editor.org/info/rfc5881>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

9.2. Informative References

- [BBF-TR-146] Broadband Forum, "BBF Technical Report - Subscriber Sessions Issue 1", 2013, <<https://www.broadband-forum.org/technical/download/TR-146.pdf>>.
- [I-D.wang-bfd-one-arm-use-case] Wang, R., Cheng, W., Zhao, Y., and A. Liu, "Using One-Arm BFD in Cloud Network", Work in Progress, Internet-Draft, draft-wang-bfd-one-arm-use-case-00, 18 November 2019, <<https://datatracker.ietf.org/doc/html/draft-wang-bfd-one-arm-use-case-00>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/

RFC3704, March 2004, <<https://www.rfc-editor.org/info/rfc3704>>.

[RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <<https://www.rfc-editor.org/info/rfc5082>>.

[RFC8704] Sriram, K., Montgomery, D., and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Forwarding", BCP 84, RFC 8704, DOI 10.17487/RFC8704, February 2020, <<https://www.rfc-editor.org/info/rfc8704>>.

Authors' Addresses

Weiqiang Cheng
China Mobile
Beijing
China

Email: chengweiqiang@chinamobile.com

Ruixue Wang
China Mobile
Beijing
China

Email: wangruixue@chinamobile.com

Xiao Min (editor)
ZTE Corp.
Nanjing
China

Email: xiao.min2@zte.com.cn

Reshad Rahman
Equinix
Ottawa
Canada

Email: reshad@yahoo.com

Raj Chetan Boddireddy
Juniper Networks

Email: rchetan@juniper.net